arXiv:1511.07245v1 [math.LO] 23 Nov 2015

CHARACTERIZING MODEL-THEORETIC DIVIDING LINES VIA COLLAPSE OF GENERALIZED INDISCERNIBLES

VINCENT GUINGONA, CAMERON DONNAY HILL, AND LYNN SCOW

Abstract. We use the notion of collapse of generalized indiscernible sequences to classify various model theoretic dividing lines. In particular, we use collapse of n-multi-order indiscernibles to characterize op-dimension n; collapse of function-space indiscernibles (i.e. parameterized equivalence relations) to characterize rosy theories; and finally, convex equivalence relation indiscernibles to characterize NTP2 theories.

1. Introduction

In model theory, and in S. Shelah's classification theory in particular, one central program is the search for robust dividing lines among complete theories – dividing lines between intelligibility and non-structure. For a dividing line to be sufficiently interesting, one often desires that both sides of the line have interesting mathematical content. Moreover, if such a dividing line has multiple characterizations coming from seemingly different contexts, this lends credence to the notion that the line is indeed important. The exemplar of this in model theory is the notion of stability, first introduced by S. Shelah [14]. Both stable and unstable theories are inherently interesting, and stability enjoys many different characterizations, from cardinalities of Stone spaces, to coding orders, to collapse of indiscernible sequences to indiscernible sets. Another example of such a robust dividing line is NIP; like stability, there are important structures with theories on both sides of the line, and the model theory on both sides can be very rich.

In this paper, we focus on the notion of collapsing generalized indiscernibles as a means of characterizating/defining various model-theoretic dividing lines that are already well-established in the literature. It can be said that this work started with Shelah's original characterization of stable theories, Theorem II.2.13 of [14]. There he shows that a theory is stable if and only if every indiscernible sequence is an indiscernible set (i.e., the order of the sequence "does not matter" or is invisible to models of the theory). In a similar fashion, it can be shown that other well-known model-theoretic dividing lines can be characterized similarly by such "collapse" statements, which we formalize here. In order to carry this out, we must expand our definition of "indiscernibility." A quintessential example of this phenomenon beyond stability is the third author's characterization of NIP theories by ordered-graph indiscernibles. In [10], she shows that a theory is NIP if and only if any indiscernible "picture" of the generic ordered graph is actually indiscernible *without* the graph structure (i.e., it collapses to order in the sense the the graph relation is invisible to models of the theory).

The project in this paper was also alluded to in the work of the first and second authors on op-dimension [3], and this will be discussed further in Section 3 below. Generalized indiscernibles were first introduced by Shelah (see Section VII.2 of [14]), where they were used in the context of tree-indexed indiscernibles in the hope of understanding the tree property (see, for example, Theorem III.7.11 of [14]), but propose a slightly different formulation suggested by the second author in presently-unpublished work – see [4] – which slightly simplifies some aspects of Section 4. For more on tree indiscernibles in particular, see [6, 7].

Outline of the paper. In Section 2, we introduce all the relevant notation and notions for generalized indiscernibility, the Ramsey Property, and the Modeling Property. As noted above, this will mostly follow [4]. In the remaining sections, we will apply this to specific cases, exhibiting the "collapse" characterizations of various dividing lines. In Section 3, we begin with the example of n-multi-order indiscernibles, first considered in [3], and show that collapse down to n orders characterizes op-dimension n (Theorem 3.4 below). In Section 4, we consider function-space indiscernibles, and show that collapse down to two indiscernible sequences characterizes rosiness (Theorem 4.7 below). In Section 5, we consider convex equivalence relation indiscernibles, showing that NTP2 is equivalent to a dichotomy between collapsing down to an indiscernible sequence or having dividing witnessed by a particular formula (Theorem 5.8 below).

2. Theories of (generalized) indiscernibles

In this section, we review the definitions associated with structural Ramsey theory: Fraïssé-like theories, the Ramsey Property, and the Modeling Property, and theories of indiscernibles. The development (taking up subsections 2.0.1 through 2.2) is just a summary of that in [4], which in turn is based in large part on [10]. In Subsection 2.4, we formalize the notion of "collapse of indiscernibles" as it will be used in this article.

2.0.1. *Notation and conventions for finite structures.* Mirroring the notation used in [4], we adopt a somewhat eccentric streamlined notation for finite structures.

Definition 2.1 (Eccentric notation for finite structures)**.** In order to save ourselves from writing "...is a finite structure" and "is a finite substructure of" *ad nauseum*, we establish a slightly eccentric convention. With the one exception of ℓ-embeddings (see the next subsection), upper-case calligraphic letters – $\mathcal{A}$, $\mathcal{B}$ and so forth – always denote *infinite* structures. On the other hand, lower-case gothic letters – $\mathfrak{a}, \mathfrak{b}, \mathfrak{c}$ and so forth – always denote *finite* structures. Thus, $\mathfrak{a} \leq \mathcal{A}$, for example, means that $\mathfrak{a}$ is a finite (induced) substructure of $\mathcal{A}$. Moreover, given a 1-sorted structure $\mathfrak{a}$, $\|\mathfrak{a}\|$ is its universe (a finite set), and $|\mathfrak{a}|$ is the cardinality of $\mathfrak{a}$.

We write $\mathrm{Emb}(\mathfrak{a}, \mathcal{A})$, $\mathrm{Emb}(\mathfrak{a}, \mathfrak{b})$, $\mathrm{Emb}(\mathcal{A}, \mathcal{B})$ for the sets of all embeddings $\mathfrak{a} \to \mathcal{A}$, $\mathfrak{a} \to \mathfrak{b}$ and $\mathcal{A} \to \mathcal{B}$, respectively. For an embedding u in one of these sets, $u\mathfrak{a}$ (or $u\mathcal{A}$) is the substructure of the codomain induced on the image of u.

Remark 2.2 (Language Convention). Although the machinery of theories of indiscernibles accomodates countably infinite relational languages (even allowing finitely many constant symbols), in this article, the signature $\mathrm{sig}(\mathscr{L}_0)$ of every language of a Fraïssé-like theory $(T_0, \mathscr{F}_0)$ will be finite relational.

2.1. **Fraïssé-like theories.** Since our languages will never have function symbols, it is necessary to accommodate "base theories" that are not universal. To do this, we define the notion of a fragment – a surrogate for the quantifier-free formulas – and the notion of a Fraïssé-like theory, which is analogous to the universal theory of a genuine Fraïssé class. A theory in a language with a function symbol that is known to be a theory of indiscernibles can be converted to a purely relational theory over a base theory where the base theory states that the new relation symbol behaves as the old function symbol did.

Definition 2.3 (Fragments)**.** Given a language $\mathscr{L}_0$, we define a *fragment* of $\mathscr{L}_0$ to be a subset $\mathscr{F}$ of $\mathscr{L}_0$-formulas with the following closure properties:

- $\mathscr{F}$ contains every quantifier-free $\mathscr{L}_0$-formula.
- $\mathscr{F}$ is closed under taking boolean combinations and under changes of free variables.
- $\mathscr{F}$ is closed under taking subformulas.

We write $\forall(\mathscr{F})$ for the set of formulas of the form $\forall\overline{x}\,\varphi$ where $\varphi \in \mathscr{F}$. Abusing the notation somewhat, we also write $\forall(\mathscr{F})$ for the closure of this set under conjunctions and disjunctions.

We say that $\mathscr{F}$ is the *trivial fragment* if it is the set of quantifier-free formulas of $\mathscr{L}_0$.

Definition 2.4 (Fraïssé-like theories)**.** Given a language $\mathscr{L}_0$, a *Fraïssé-like theory* is a pair $(T_0, \mathscr{F})$, where $\mathscr{F}$ is a fragment of $\mathscr{L}_0$ and T_0 is a set of $\forall(\mathscr{F})$-sentences, satisfying the following conditions:

(JEP) For all $\mathfrak{a}_0, \mathfrak{a}_1 \vDash T_0$, there is a $\mathfrak{b} \vDash T_0$ such that $\mathrm{Emb}_{\mathscr{F}}(\mathfrak{a}_i, \mathfrak{b}) \neq \emptyset$ $(i < 2)$.

(AP) For $\mathfrak{a}, \mathfrak{b}_0, \mathfrak{b}_1 \vDash T_0$ and $f_i \in \mathrm{Emb}_{\mathscr{F}}(\mathfrak{a}, \mathfrak{b}_i)$ $(i < 2)$, there are $\mathfrak{c} \vDash T_0$ and $g_i \in \mathrm{Emb}_{\mathscr{F}}(\mathfrak{b}_i, \mathfrak{c})$ $(i < 2)$ such that $g_1 f_1 = g_0 f_0$.

(LS) There is a function $\lambda : \omega \to \omega$ such that for any $\mathcal{B} \vDash T_0$ and $X \subseteq_{\mathrm{fin}} B$, there is an $\mathfrak{a} \preceq_{\mathscr{F}} \mathcal{B}$ such that $X \subseteq \|\mathfrak{a}\|$ and $|\mathfrak{a}| \leq \lambda(|X|)$.

As with Fraïssé classes and Fraïssé-limits, Fraïssé-like theories support countably infinite generic models. The following proposition summarizes this fact.

Proposition 2.5. *Let $(T_0, \mathscr{F})$ be a Fraïssé-like theory in $\mathscr{L}_0$. Uniquely up to isomorphism, there is a countable $\mathscr{L}_0$-structure $\mathcal{A}$ such that:*

- *For all $\mathfrak{a} \vDash T_0$, $\mathrm{Emb}_{\mathscr{F}}(\mathfrak{a}, \mathcal{A}) \neq \emptyset$.*
- *For any $\mathfrak{a}, \mathfrak{b} \vDash T_0$ with $\mathfrak{a} \preceq_{\mathscr{F}} \mathfrak{b}$, for any $f_0 \in \mathrm{Emb}_{\mathscr{F}}(\mathfrak{a}, \mathcal{A})$, there is an $f \in \mathrm{Emb}_{\mathscr{F}}(\mathfrak{b}, \mathcal{A})$ such that $f_0 \subseteq f$.*

Moreover, $T_0^* := Th(\mathcal{A})$ is $\aleph_0$-categorical and model-complete relative to $\mathscr{F}$.

Definition 2.6 ($\mathscr{F}$-model-completion; generic models)**.** In the notation of Proposition 2.5, For a Fraïssé-like theory $(T_0, \mathscr{F})$, T_0^* is the *$\mathscr{F}$-model-completion of T_0*.

The structure $\mathcal{A} \vDash T_0^*$ is called the *generic model* of $(T_0, \mathscr{F})$.

Remark 2.7. In the case that $\mathscr{F}$ is the trivial fragment, T_0^* is the usual model-completion of T_0.

Here we give a list of all Fraïssé-like theories $(T_0, \mathscr{F})$ discussed in the paper.

Definition 2.8 (Index of theories)**.**

- T_0: stands for an arbitrary index theory
- T_n, $0 < n < \omega$:
  - Let $\mathscr{L}_n$ be the n-sorted language with sorts $\mathsf{X}_0, ..., \mathsf{X}_{n-1}$ and for each $i < n$, a binary relation symbol $<_i$ on X_i. Let T_n be the $\mathscr{L}_n$-theory asserting that each $<_i$ is a linear ordering of X_i.
  - Let $\mathscr{F}_n$ be the trivial fragment.
- T_1: this is the theory of linear order (see above)
- $T_{n-\mathrm{mlo}}$, $0 < n \leq \omega$:
  - Let $\mathscr{L}_{n-\mathrm{mlo}} = \{<_0, ..., <_{n-1}\}$ be the 1-sorted language with n binary relations. Let $T_{n-\mathrm{mlo}}$ be the $\mathscr{L}_{n-\mathrm{mlo}}$-theory which asserts that $<_i$ is a linear order for each $i < n$.
  - Let $\mathscr{F}_{n-\mathrm{mlo}}$ to be the trivial fragment.
- T_{fn}, $T_{\mathrm{eq:n}}$, $0 < n \leq \omega$:
  - Let $\mathscr{L}_{\mathrm{fn}}$ be the two-sorted language with sorts X and F, two binary relation symbols $<_{\mathsf{X}}$ and $<_{\mathsf{F}}$ on X and F, respectively, and one more relation symbol $E \subseteq \mathsf{X} \times \mathsf{F} \times \mathsf{F}$. Define T_{fn} to contain the following:
    - $(\forall x \in \mathsf{X})$ "$E_x = E(x, \cdot, \cdot)$ is an equivalence relation on F."
    - $(\forall u, u' \in \mathsf{F})\big(u = u' \leftrightarrow (\forall x \in \mathsf{X})E(x, u, u')\big)$.
    - "$<_{\mathsf{X}}$ is a linear order of X" and "$<_{\mathsf{F}}$ is a linear order of F."
    - For all n,

$(\forall x_0 <_{\mathsf{X}} \cdots <_{\mathsf{X}} x_{n-1} <_{\mathsf{X}} x_n \in \mathsf{X})(\forall u \in \mathsf{F})$"$u/\cap_{i \leq n} E_{x_i}$ is an $<_{\mathsf{F}}$-convex subset of $u/\cap_{i<n} E_{x_i}$"

– Let $\mathscr{F}_{\text{fn}}$ be the smallest fragment of $\mathscr{L}_{\text{fn}}$ containing the formula $(\exists x \in X)\neg E(x, u, u')$.
– For each $0 < n < \omega$, let $T_{\text{eq}:n}$ be the the $\mathscr{F}_{\text{fn}}$-universal theory obtained from T_{fn} by adding the sentence

$$(\forall x_0, ..., x_{n-1}, x_n \in \mathsf{X}) \bigvee_{i<n} x_n = x_i.$$

- T_{ceq}:
  – Let $\mathscr{L}_{\text{ceq}} = \{E, <\}$ be a language with two binary relations. Define the theory of *convexly-ordered equivalence relations*, T_{ceq}, to be the theory in $\mathscr{L}_{\text{ceq}}$ given by defining E to be an equivalence relation and $<$ to be a linear ordering such that equivalence classes are convexly ordered, i.e. the following holds in any model of T_{ceq}:

$$a < b < c \wedge E(a, c) \Rightarrow E(a, b) \wedge E(b, c)$$

  – Let $\mathscr{F}$ denote the trivial fragment.

To conclude this subsection, we introduce the notion of an indiscernible picture of a model of T_0^* for a Fraïssé-like theory $(T_0, \mathscr{F})$. Of course, not all Fraïssé-like theories admit indiscernible pictures, but the definition is sensible even so. By $\text{tp}^{\mathcal{A}}_{\mathscr{F}}(\overline{a})$ we mean the formulas in the fragment $\mathscr{F}$ satisfied by $\overline{a}$ in $\mathcal{A}$.

Definition 2.9 (Indiscernible pictures)**.** Let $(T_0, \mathscr{F})$ be a Fraïssé-like theory, and let $\mathcal{A} \models T_0^*$ (possibly uncountable). Let $\mathcal{M}$ be an infinite structure. An *indiscernible picture* $I : \mathcal{A} \to \mathcal{M}$ (of $\mathcal{A}$) is an injection $I : A \to M$ such that for all $0 < n < \omega$ and $\overline{a}, \overline{b} \in A^n$,

$$\text{tp}^{\mathcal{A}}_{\mathscr{F}}(\overline{a}) = \text{tp}^{\mathcal{A}}_{\mathscr{F}}(\overline{b}) \;\Rightarrow\; \text{tp}^{\mathcal{M}}(I\overline{a}) = \text{tp}^{\mathcal{M}}(I\overline{b}).$$

By an indiscernible T_0-picture in $\mathcal{M}$, we mean an indiscernible picture I of some $\mathcal{A} \models T_0^*$.

2.2. **Ramsey and Modeling Properties.** In this subsection, we establish the twin properties that identify generalized indiscernibles (theories of indiscernibles) among all Fraïssé-like theories.

The first of these – the Finitary Ramsey Property – is precisely the hypothesis that allows the classical proof of the existence of indiscernible sequences (via Ramsey Theorem and compactness) to go through largely unchanged.

Definition 2.10 (Finitary Ramsey Property)**.** Let $(T_0, \mathscr{F})$ be a Fraïssé-like theory.

- For $\mathfrak{a} \models T_0$, we say that $(T_0, \mathscr{F})$ has the $\mathfrak{a}$-Ramsey Property if for all $2 \leq k < \omega$ and all $\mathfrak{b} \models T_0$, there is a $\mathfrak{c} \models T_0$ such that for any coloring $\xi : \text{Emb}(\mathfrak{a}, \mathfrak{c}) \to k$, there is an embedding $u \in \text{Emb}(\mathfrak{b}, \mathfrak{c})$ such that for all $f_0, f_1 \in \text{Emb}(\mathfrak{a}, \mathfrak{b})$, $\xi(u \circ f_0) = \xi(u \circ f_1)$. (This condition is denoted by the Rado partition arrow: $\mathfrak{c} \to (\mathfrak{b})^{\mathfrak{a}}_k$.)
- $(T_0, \mathscr{F})$ has the Ramsey Property if it has the $\mathfrak{a}$-Ramsey Property for every $\mathfrak{a} \models T_0$.

The second batch of definitions is more clearly identifiable as characterizing a "theory of indiscernibles." The reader may find it helpful to recall that for any infinite structure $\mathcal{M}$ and any sequence $(a_i)_{i<\omega}$, there is an indiscernible sequence $(a'_i)_{i<\omega}$ (in an elementary extension of $\mathcal{M}$) "patterned" on $(a_i)_{i<\omega}$. The first two definitions below (of ℓ-embeddings and EM-templates) formalize the idea of patterning. The Modeling Property, then, uses these notions to describe what is, naturally, a "theory of indiscernibles."

Definition 2.11 (ℓ-Embeddings)**.** Let $(T_0, \mathscr{F})$ be a Fraïssé-like theory, and let $\mathcal{A}, \mathcal{B} \models T_0^*$. By an "$\ell$-embedding $f : \mathcal{A} \rightsquigarrow \mathcal{B}$," we mean a family $f = \big(f_X : \mathcal{A}_X \to \mathcal{B}\big)_{X \subset_{\text{fin}} A}$ of embeddings, where for each $X \subset_{\text{fin}} A$, $\mathcal{A}_X \prec_{\mathscr{F}} \mathcal{A}$ is a finite model of T_0 containing X. Implicitly, there is a uniform bounding function $s : \omega \to \omega$ such that $|\mathcal{A}_X| \leq s(|X|)$ for all $X \subset_{\text{fin}} A$.

Definition 2.12 (EM-templates; patterning)**.** Let $(T_0, \mathscr{F})$ be a Fraïssé-like theory, and let $\mathcal{A} \vDash T_0^*$. In an arbitrary language $\mathscr{L}$, let $\mathcal{M}$ be some $\mathscr{L}$-structure. We define an *EM-template* to be a family Γ of maps $A^n \to M^n : \overline{a} \mapsto \Gamma(\overline{a})$ $(0 < n < \omega)$ satisfying the following:

- For all $0 < n < \omega$, $\overline{a} \in A^n$, $\overline{b} = \Gamma(\overline{a})$, and $\sigma \in Sym(n)$,
$$\Gamma(a_{\sigma(0)}, ..., a_{\sigma(n-1)}) = (b_{\sigma(0)}, ..., b_{\sigma(n-1)}).$$
- For all $0 < m, n < \omega$, $\overline{a} \in A^m$ and $\overline{a}' \in A^n$,
$$\mathrm{tp}^{\mathcal{M}}(\Gamma(\overline{a}^\frown \overline{a}')) = \mathrm{tp}^{\mathcal{M}}(\Gamma(\overline{a})^\frown \Gamma(\overline{a}'))$$

Now, suppose $\mathcal{M} \preceq \mathcal{N}$, $\mathcal{B} \vDash T_0^*$, $I : \mathcal{B} \to \mathcal{N}$ is an indiscernible picture of $\mathcal{B}$ in $\mathcal{N}$. We say that I *is patterned on* Γ if there is a family $\left(f^\Delta : \mathcal{B} \rightsquigarrow \mathcal{A}\right)_{\Delta \subset_{\mathrm{fin}} \mathscr{L}}$ of patterning maps such that for all $\Delta \subset_{\mathrm{fin}} \mathscr{L}$, $X \subset_{\mathrm{fin}} B$ and $\overline{b} \in X^{<\omega}$, $\mathrm{tp}^{\mathcal{N}}_\Delta(I\overline{b}) = \mathrm{tp}^{\mathcal{M}}_\Delta\left(\Gamma(f^\Delta_X \overline{b})\right)$.

Definition 2.13 (Modeling Property)**.** A Fraïssé-like theory $(T_0, \mathscr{F})$ with generic model $\mathcal{A} \vDash T_0^*$ is said to have the *Modeling Property* if for every infinite structure $\mathcal{M}$ and every EM-template $\Gamma : A^{<\omega} \to M^{<\omega}$, there are $\mathcal{M} \preceq \mathcal{N}$ and an indiscernible picture $I : \mathcal{A} \to \mathcal{N}$ patterned on Γ.

Finally, we can state the main characterization theorem connecting the Ramsey Property and the Modeling Property. A very-slightly weaker form of this theorem was proved in [10], and in [4], it was proved as one of a somewhat longer list of equivalent conditions.

Theorem 2.14. *Let $(T_0, \mathscr{F})$ be a Fraïssé-like theory. Then $(T_0, \mathscr{F})$ has the Finitary Ramsey Property if and only if it has the Modeling Property.*

Definition 2.15 (Theory of indiscernibles)**.** A *theory of indiscernibles* is precisely a Fraïssé-like theory $(T_0, \mathscr{F})$ that has the Modeling Property (equivalently, the Finitary Ramsey Property).

2.3. **Examples.** We hope that the following examples will be helpful to the reader in understanding the definitions that we have just laid down.

Example 2.16. $(T_n, \mathscr{F}_n)$ is a Fraïssé-like theory, and T_n^* is the theory asserting just that each $(\mathsf{X}_i, <_i)$ is a dense linear order without endpoints. If $\mathcal{A} \vDash T_n^*$ is countable, then $Aut(\mathcal{A}) = Aut(\mathbb{Q}, <)^n$, which is extremely amenable. So by Lemma 6.7 of [5] we can conclude that $(T_n, \mathscr{F}_n)$ is a theory of indiscernibles.

Now, consider an indiscernible picture $I : \mathcal{A} \to \mathcal{M}$, where $\mathcal{M}$ is some infinite $\mathscr{L}$-structure. We note that for $i < j < n$ and $a \in \mathsf{X}_i^{\mathcal{A}}$, $b \in \mathsf{X}_i^{\mathcal{A}}$, $I(a)$ and $I(b)$ may be of different sorts. However, each sequence $I_j = (a_t)_{t \in \mathsf{X}_i^{\mathcal{A}}}$ $(j < n)$ is an indiscernible sequence in the usual sense, and moreover, the family $(I_0, ..., I_{n-1})$ comprises mutually indiscernible sequences in the sense that for each $j < n$, I_j is indiscernible over $\bigcup_{k<n:k\neq j} I_k$.

Example 2.17 (n-multi-order). $(T_{n-\mathrm{mlo}}, \mathscr{F}_{n-\mathrm{mlo}})$ is a Fraïssé-like theory, and $T^*_{n-\mathrm{mlo}}$ is the theory where each $<_i$ is a dense linear order and these orders are independent. Then $T^*_{n-\mathrm{mlo}}$ is $\aleph_0$-categorical and admits elimination of quantifiers. Moreover, $T_{n-\mathrm{mlo}}$ is a theory of indiscernibles. For more details, see Definition 1.2 of [3], where this is called MLO^*_n. Also see Corollary 1.4 of [1] and Theorem 4 of [15].

Example 2.18 (function spaces). T_{fn} is a $\mathscr{F}_{\mathrm{fn}}$-universal theory. It is easy to see that $(T_{\mathrm{fn}}, \mathscr{L}_{\mathrm{fn}})$ is a Fraïssé-like theory. In fact, the model companion shares an automorphism group with that of the Shelah tree $I_s = ({}^{\omega>}\omega, \lhd, \wedge, <_{\mathrm{lex}}, \{P_n\}_n)$ from [7] where $\wedge$ is the meet in the partial tree order $\lhd$, $<_{\mathrm{lex}}$ is the lexicographical order on sequences, and P_n is a unary predicate picking out the n-th level of the tree. Identify leaves of ${}^{n\geq}\omega$ with the one-point classes $[x]/\bigcap_{1\leq 1\leq n} E_{x_i}$.

Example 2.19. $(T_{\mathrm{ceq}}, \mathscr{F})$ is a Fraïssé-like theory. A countable model of the model companion T^*_{ceq} is a collection of densely ordered convex equivalence classes, each of which is a dense linear order.

2.4. **Collapse.** It probably is not surprising (given an understanding of stable theories) that for a given theory T, not every indiscernible picture of a model of T_0^* "survives" – sometimes, the indiscernible picture is already indiscernible relative to a reduct of $(T_0, \mathscr{F})$. This phenomenon is the basis for our approach to dividing lines in this article, so obviously, we must formalize exactly what we mean.

Definition 2.20 (Reducts of Fraïssé-like theories)**.** Let $(T_0, \mathscr{F}_0)$ and $(T_0^+, \mathscr{F}_0^+)$ be Fraïssé-like theories in $\mathscr{L}_0$ and $\mathscr{L}_0^+$, respectively, and let $\mathcal{A}_0 \vDash T_0^*$ and $\mathcal{A}_0^+ \vDash (T_0^+)^*$ be their respective generic models.

- As usual, we say that $\mathcal{A}_0$ is a reduct of $\mathcal{A}_0^+$, if there is a bijection $w : A_0 \to A_0^+$ such that for every formula $\varphi(x_0, ..., x_{n-1})$ of $\mathscr{L}_0$, there is a formula $\psi(x_0, ..., x_{n-1})$ of $\mathscr{L}_0^+$ such that
$$\mathcal{A}_0 \vDash \varphi(\overline{a}) \Leftrightarrow \mathcal{A}_0^+ \vDash \psi(w\overline{a})$$
for all $\overline{a} \in A_0^n$. As all of the theories in question are $\aleph_0$-categorical, this is equivalent to asserting that, up to the identification provided by w, every complete type of T_0^* over $\emptyset$ is equal to a (finite) union of complete types of $(T_0^+)^*$ over $\emptyset$. We will say that $(T_0, \mathscr{F}_0)$ is a reduct of $(T_0^+, \mathscr{F}_0^+)$ if the generic model of the former is a reduct of the generic model of the latter.

Definition 2.21 (Reducts of Fraïssé-like theories)**.** Let $(T_0, \mathscr{F}_0)$ and $(T_0^+, \mathscr{F}_0^+)$ be theories of indiscernibles, and let T be an arbitrary complete theory. Assume $(T_0, \mathscr{F}_0)$ is a reduct of $(T_0^+, \mathscr{F}_0^+)$ – so fix $w : A_0 \to A_0^+$ witnessing this, where $\mathcal{A}_0$ and $\mathcal{A}_0^+$ are the generic models of the two theories, respectively.

- Let $I : \mathcal{A}_0^+ \to \mathcal{M}$ be an indiscernible picture of $\mathcal{A}_0^+ \vDash (T_0^+)^*$ in a model $\mathcal{M}$ of T. We say that I *collapses to* $(T_0, \mathscr{F}_0)$ if the composite $I \circ w : \mathcal{A}_0 \to \mathcal{M}$ is an indiscernible picture of $\mathcal{A}_0$ in $\mathcal{M}$.
- We say that T *collapses* $(T_0^+, \mathscr{F}_0^+)$ *to* $(T_0, \mathscr{F}_0)$ if every indiscernible picture $I : \mathcal{A}_0^+ \to \mathcal{M}$ of $\mathcal{A}_0^+ \vDash (T_0^+)^*$ in a model $\mathcal{M}$ of T collapse to $(T_0, \mathscr{F}_0)$.

Remark 2.22. We note that our formalization of collapse requires that a specific reduct is isolated before hand. Although it is also possible (presumably) to speak of "collapse" without requiring that the remainder is anything in particular, this is *not* the approach we adopt here.

2.5. **Dividing across.** Although we can get similar collapse results for op-dimension and rosiness, for NTP2, we will need to consider a slightly different notion. This is primarily because of the fact that, while stability, NIP, op-dimension, and rosiness are characterized by the non-existence of a pattern with only existential criteria, NTP2 (and likewise, NTP1 for that matter) is characterized by the non-existence of a pattern with both existential and universal requirements. So we will need a notion of dividing across a fixed sequence I to deal with the "universal part" of the pattern.

Definition 2.23 (Dividing across I)**.** Let $I : \mathcal{A} \to \mathcal{M}^y$ be an indiscernible picture of $\mathcal{A} \vDash T_0^*$, and let $\varphi(x; y)$ be an $\mathscr{L}$-formula. We say that $\varphi(x; y)$ *divides across* I if there is an indiscernible sequence $\langle a_i \rangle_{i<\omega}$ in $\mathcal{A}$ such that:
$$\{\varphi(x; Ia_i)\}_{i<\omega} \text{ is inconsistent.}$$

Notice that saying $\varphi(x; y)$ divides across I, witnessed by $\langle a_i \rangle_{i<\omega}$, shows that $\varphi(x; Ia_0)$ divides over $\emptyset$, since $\langle Ia_i \rangle_{i<\omega}$ is an indiscernible sequence in $\mathcal{M}$. As usual, by compactness, this is equivalent to demanding that $\{\varphi(x; Ia_i)\}_{i<\omega}$ is k-inconsistent for some $k < \omega$.

We will use this notion in Section 5 below. We will introduce a specialized versions of dividing, namely "dividing vertically." To connect diving to indiscernible collapse, we will also have a notion of *coding* formulas. Though this definition depends on the specific index theory T_0, the basic idea

is to have a single $\mathscr{L}$-formula witnessing the collapse of a particular relation symbol in $\mathscr{L}_0$ for a indiscernible picture I.

3. *op*-Dimension and indiscernible multi-orders

Let T be a complete first-order $\mathscr{L}$-theory and let $\mathbb{U}$ be a monster model for T.

3.1. **Background on op-dimension.** There are many equivalent definitions for op-dimension given in [3]. The main one of interest here will be the following.

Definition 3.1 (op-dimension via IRD-patterns)**.** We say that a partial type $\pi(x)$ has *op-dimension* $\geq n$ if there exists formulas $\psi_i(x; y_i)$ and sequences $\langle b_{i,j} : j < \omega\rangle$ for $i < n$ such that, for all functions $f : n \to \omega$, the following partial type is consistent

$$\pi(x) \cup \{\psi_i(x; b_{i,j}) : j < f(i), i < n\} \cup \{\neg\psi_i(x; b_{i,j}) : j \geq f(i), i < n\}.$$

We will call such $\langle \psi_i, \langle b_{i,j} : j < \omega\rangle : i < n\rangle$ an *IRD-pattern* of length n and depth ω in $\pi(x)$. We will also denote this by $\mathsf{opD}(\pi) \geq n$. We will write $\mathsf{opD}(a/B)$ for $\mathsf{opD}(\mathrm{tp}(a/B))$. If the op-dimension of π is $\geq n$ for all n, we say it is infinite and write $\mathsf{opD}(\pi) = \infty$.

Notice that by the modeling property in Example 2.16, we may assume that the $b_{i,j}$ in the definition of op-dimension are such that $\{\langle b_{i,j} : j \in \mathbb{Q}\rangle : i < n\}$ is mutually indiscernible (and indexed by $\mathbb{Q}$ or, in fact, any infinite order J).

Notice that if $x = x$ has finite op-dimension, then in particular T has NIP in the sort of x. The converse does not hold. Consider $T^*_{\omega-\mathrm{mlo}}$ (see Definition 2.8). A partial type π is stable if and only if $\mathsf{opD}(\pi) = 0$. It is easy to check that the op-dimension of a partial type π is bounded above by the dp-rank of π (any IRD-pattern can be converted to an ICT-pattern). In general, however, op-dimension can be much smaller (e.g., there are stable, non-strongly dependent theories). On the other hand, op-dimension enjoys many of the same properties as dp-rank, including subadditivity (see Theorem 2.2 of [3]). That is,

$$\mathsf{opD}(ab/A) \leq \mathsf{opD}(a/A) + \mathsf{opD}(b/Aa).$$

If T is o-minimal (or simply distal), then dp-rank and op-dimension coincide (Theorem 3.1 of [3]).

Theorem 3.2 (Theorem 1.21 and Proposition 2.1 of [3])**.** *For a partial type $\pi(x)$ over a set B, the following are equivalent*

1. *$\mathsf{opD}(\pi) \leq n$,*
2. *For all formulas $\varphi(x; y)$, all sequences $\langle b_i : i \in \mathbb{Q}\rangle$ indiscernible over B, and for all $a \vDash \pi$, there exists a partition of $\mathbb{Q}$ into convex sets $C_0 < ... < C_n$ such that, for all $\ell \leq n$, the set $\{i \in C_\ell : \vDash \varphi(a; b_i)\}$ is either finite or cofinite in C_ℓ.*
3. *For all $\{\langle b_{i,j} : j \in J\rangle : i \leq n\}$ almost mutually indiscernible over B and all $a \vDash \pi$, there exists $i \leq n$ such that $\langle b_{i,j} : j \in J\rangle$ is almost indiscernible over Ba.*

3.2. **n-Multi-order characterization of op-dimension.**

Proposition 3.3 (Proposition 1.18 of [3])**.** *For a partial type $\pi(x)$ and $n < \omega$, the following are equivalent:*

1. *$\mathsf{opD}(\pi) \geq n$,*
2. *There exists an indiscernible $T_{n-\mathrm{mlo}}$-picture $I : \mathcal{A} \to \pi(\mathbb{U})$ and a formula $\varphi(x; y)$ such that, for all $i < n$ and for all $X \subseteq A$ a cut of $(A; <_i)$, there exists $b \in \mathbb{U}$ such that $X = \{a \in \mathcal{A} : \vDash \varphi(Ia; b)\}$.*

Along similar lines, we can use collapse of indiscernible $T_{n-\mathrm{mlo}}$-pictures to characterize op-dimension, supposing the ambient theory has NIP. The following is the main result of this section.

Theorem 3.4 (Characterization of op-dimension by collapsing indiscernibles)**.** *Suppose T has NIP. For a partial type $\pi(x)$ and $n < \omega$, the following are equivalent:*

(1) $\mathsf{opD}(\pi) < n$,
(2) *For all indiscernible $T_{n-\mathrm{mlo}}$-pictures $I : \mathcal{A} \to \pi(\mathbb{U})$, there exists $j < n$ such that I is an indiscernible picture for $\mathscr{F}^{-j}_{n-\mathrm{mlo}} := \{<_i : i \neq j, i < n\}$*

Proof. (1) $\Rightarrow$ (2): Suppose (2) fails. That is, there exists an indiscernible $T_{n-\mathrm{mlo}}$-picture $I : \mathcal{A} \to \pi(\mathbb{U})$ such that, for all $j < n$, I is not an indiscernible picture for $\mathscr{F}^{-j}_{n-\mathrm{mlo}} := \{<_i : i \neq j, i < n\}$. That is, there exists $k < \omega$ and $\overline{a}^j, \overline{b}^j \in A^k$ such that

$$\mathrm{tp}^{\mathcal{A}}_{\mathscr{F}^{-j}_{n-\mathrm{mlo}}}(\overline{a}^j) = \mathrm{tp}^{\mathcal{A}}_{\mathscr{F}^{-j}_{n-\mathrm{mlo}}}(\overline{b}^j), \mathrm{tp}^{\mathcal{A}}_{\mathscr{F}_{n-\mathrm{mlo}}}(\overline{a}^j) \neq \mathrm{tp}^{\mathcal{A}}_{\mathscr{F}_{n-\mathrm{mlo}}}(\overline{b}^j), \text{ and } \mathrm{tp}^{\mathbb{U}}_{\mathscr{L}}(I\overline{a}^j) \neq \mathrm{tp}^{\mathbb{U}}_{\mathscr{L}}(I\overline{b}^j).$$

We may assume the same k works for all $j < n$ (take the maximal such) and the coordinates of $\overline{a}^j$ are distinct and the coordinates of $\overline{b}^j$ are distinct. Without loss of generality, by reordering, suppose

$$a^j_0 <_j a^j_1 <_j \ldots <_j a^j_{k-1}.$$

Since transpositions generate the symmetric group on k, we may assume that the $<_j$-order type of $\overline{a}^j$ and $\overline{b}^j$ differ by transposing two elements. That is, for some $t_j < k-1$,

$$b^j_0 <_j \ldots <_j b^j_{t_j-1} <_j b^j_{t_j+1} <_j b^j_{t_j} <_j b^j_{t_j+2} <_j \ldots <_j b^j_{k-1}.$$

Since $\mathrm{tp}^{\mathbb{U}}_{\mathscr{L}}(I\overline{a}^j) \neq \mathrm{tp}^{\mathbb{U}}_{\mathscr{L}}(I\overline{b}^j)$, there exists $\psi_j(y_0, \ldots, y_{k-1}) \in \mathscr{L}$ such that

$$\models \psi_j(I\overline{a}^j) \wedge \neg\psi_j(I\overline{b}^j).$$

Fix $m < \omega$. We use this as a template to construct an IRD-pattern of depth m and length n in π, showing that $\mathsf{opD}(\pi) \geq n$ (and hence (1) fails).

For each $j < n$, choose

$$d^j_0, \ldots, d^j_{t_j-1}, d^j_{t_j+1}, \ldots, d^j_{k-1} \in {}^n\mathbb{Q}$$

such that, for all $i < n$,

- $d^j_\ell(i) < d^j_s(i)$ if $a^j_\ell <_i a^j_s$ for all $\ell, s \neq t_j$,
- $d^j_\ell(i) < -1$ if $a^j_\ell <_i a^j_{t_j}$ for all $\ell \neq t_j$,
- $d^j_\ell(i) > m$ if $a^j_\ell >_i a^j_{t_j}$ for all $\ell \neq t_j$.

Thereby, for all $f \in {}^n(-1, m)_{\mathbb{Q}}$, for all i,

$$\langle d^j_0, \ldots, d^j_{t_j-1}, f, d^j_{t_j+1}, \ldots, d^j_{k-1}\rangle$$

has the same $<_i$-order type as $\overline{a}$. In other words, for all $f \in {}^n(-1, m)_{\mathbb{Q}}$,

$$\mathrm{tp}^{{}^n\mathbb{Q}}_{\mathscr{F}_{n-\mathrm{mlo}}}(d^j_0, \ldots, d^j_{t_j-1}, f, d^j_{t_j+1}, \ldots, d^j_{k-1}) = \mathrm{tp}^{\mathcal{A}}_{\mathscr{F}_{n-\mathrm{mlo}}}(\overline{a}^j).$$

For each $j < n$ and $\ell < m$, define $e^j_\ell \in {}^n\mathbb{Q}$ as follows: For each $i < n$,

- $e^j_\ell(i) = d^j_{t_j+1}(i)$ for all $i \neq j$,
- $e^j_\ell(j) = \ell - 1/2$.

So, for e^j_ℓ, we are merely taking $d^j_{t_j+1}$ and stretching it, in the jth coordinate, across $(-1, m+1)_{\mathbb{Q}}$. Now, for any $f \in {}^n m$, for any $j < n$ and any $\ell < m$, we have

- $d^j_0 <_j \ldots <_j d^j_{t_j-1} <_j f <_j e^j_\ell <_j d^j_{t_j+2} <_j \ldots <_j d^j_{k-1}$ for $f(j) < \ell$ and,
- $d^j_0 <_j \ldots <_j d^j_{t_j-1} <_j e^j_\ell <_j f <_j d^j_{t_j+2} <_j \ldots <_j d^j_{k-1}$ for $f(j) \geq \ell$.

In other words, for any $f \in {}^n m$, $\ell < m$, and $j < n$,

- $\mathrm{tp}^{{}^n\mathbb{Q}}_{\mathscr{F}^{-}_{n-\mathrm{mlo}}}(d^j_0, ..., d^j_{t_j-1}, f, e^j_\ell, d^j_{t_j+2}, ..., d^j_{k-1}) = \mathrm{tp}^{\mathcal{A}}_{\mathscr{F}^{-}_{n-\mathrm{mlo}}}(\overline{a}^j)$ for $f(j) < \ell$ and,
- $\mathrm{tp}^{{}^n\mathbb{Q}}_{\mathscr{F}^{-}_{n-\mathrm{mlo}}}(d^j_0, ..., d^j_{t_j-1}, f, e^j_\ell, d^j_{t_j+2}, ..., d^j_{k-1}) = \mathrm{tp}^{\mathcal{A}}_{\mathscr{F}^{-}_{n-\mathrm{mlo}}}(\overline{b}^j)$ for $f(j) \geq \ell$.

After embedding into $\mathcal{A}$ (by first extending each $<_i$ arbitrarily to be a linear order), we see that, for all $f \in {}^n m$, $\ell < m$, and $j < n$,

$$\vDash \psi_j(Id^j_0, ..., Id^j_{t_j-1}, If, Ie^j_\ell, Id^j_{t_j+2}, ..., Id^j_{k-1}) \text{ iff. } f(j) < \ell.$$

Let $\psi'_j(y; y_0, ..., y_{t_j-1}, y_{t_j+1}, ..., y_{k-1}) = \psi_j(y_0, ..., y_{t_j-1}, y, y_{t_j+1}, ..., y_{k-1})$, let $\alpha_f = If \vDash \pi$ for all $f \in {}^n m$, and let

$$\beta_{j,\ell} = \langle Id^j_0, ..., Id^j_{t_j-1}, Ie^j_\ell, Id^j_{t_j+2}, ..., Id^j_{k-1} \rangle$$

for all $j < n$ and $\ell < m$. Then, for all $j < n$, $\ell < m$, and $f \in {}^n m$,

$$\vDash \psi'_j(\alpha_c; \beta_{j,\ell}) \text{ iff. } f(j) < \ell.$$

Hence, ψ'_j and $\beta_{j,\ell}$ form an IRD-pattern of depth m and length n in π. Since m was arbitrary, by compactness, we get that $\mathsf{opD}(\pi) \geq n$, as desired.

(2) $\Rightarrow$ (1): Suppose (1) fails, so $\mathsf{opD}(\pi) \geq n$. By Proposition 3.3, there exists an indiscernible $T_{n-\mathrm{mlo}}$-picture $I : \mathcal{A} \to \pi(\mathbb{U})$ and a formula $\varphi(x; y)$ such that, for all $j < n$ and for all $X \subseteq A$ a cut of $(A; <_j)$, there exists $b \in \mathbb{U}$ such that $X = \{a \in \mathcal{A} : \vDash \varphi(Ia; b)\}$. Since T has NIP, $\varphi(x; y)$ has NIP, so fix $m < \omega$ such that φ has VC-dimension $< m$. That is,

$$\vDash \neg \exists x_0, ..., x_{m-1} \bigwedge_{s \in {}^m 2} \left(\exists y \bigwedge_{\ell < m} \varphi(x_\ell, y)^{s(\ell)} \right).$$

We show that (2) fails on I. Fix $j < n$ and suppose, by means of contradiction, that I is an indiscernible picture for $\mathscr{F}^{-j}_{n-\mathrm{mlo}} := \{<_i : i \neq j, i < n\}$. Fix any $\overline{c} \in \mathcal{A}^m$. By choice of m (i.e., that m is greater than the VC-dimension of φ), there exists $s \in {}^m 2$ such that

$$\vDash \neg \exists y \bigwedge_{\ell < m} \varphi(Ic_\ell, y)^{s(\ell)}.$$

Let $L_0 = \{\ell < m : s(\ell) = 0\}$ and $L_1 = \{\ell < m : s(\ell) = 1\}$. Choose $\overline{d} \in \mathcal{A}^m$ such that

- $\mathrm{tp}^{\mathcal{A}}_{\mathscr{F}^{-j}_{n-\mathrm{mlo}}}(\overline{c}) = \mathrm{tp}^{\mathcal{A}}_{\mathscr{F}^{-j}_{n-\mathrm{mlo}}}(\overline{d})$, and
- For all $\ell_1 \in L_1$ and $\ell_0 \in L_0$, $d_{\ell_1} <_j d_{\ell_0}$.

By indiscernibility with respect to $\mathscr{F}^{-j}_{n-\mathrm{mlo}}$,

$$\vDash \neg \exists y \bigwedge_{\ell < m} \varphi(Id_\ell, y)^{s(\ell)}.$$

Without loss of generality, by reordering, suppose $d_0 <_j d_1 <_j ... <_j d_{m-1}$, so we have

$$\vDash \neg \exists y \left(\bigwedge_{\ell < m_0} \varphi(Id_\ell; y) \wedge \bigwedge_{m_0 \leq \ell < m} \neg \varphi(Id_\ell; y) \right),$$

where $m_0 = |L_1|$. Now choose any $X \subseteq A$ a cut of $(A; <_j)$ with $d_\ell \in X$ if and only if $\ell < m_0$. Therefore, there exists $b \in \mathbb{U}$ such that $X = \{a \in \mathcal{A} : \varphi(Ia; b)\}$. In particular,

$$\vDash \bigwedge_{\ell < m_0} \varphi(Id_\ell, b) \wedge \bigwedge_{m_0 \leq \ell < m} \neg \varphi(Id_\ell, b).$$

This is a contradiction. Hence (2) fails and we are done.

□

Note that we only use NIP in the proof of Theorem 3.4, (2) $\Rightarrow$ (1). However, this use is necessary. Consider the theory T of the random graph (where $\mathscr{L} = \{R\}$) and take $\pi(x) = [x = x]$ in the homesort. Then, any indiscernible sequence is an indiscernible set (either everything is R-connected or R-disconnected). That is, Theorem 3.4 (2) holds for $n = 1$. On the other hand, Theorem 3.4 (1) fails for all n (i.e., $\mathsf{opD}(\pi) = \infty$). To see this, notice that $R(x; y)$ together with any distinct $\langle \langle b_{i,j} : j < \omega \rangle : i < n \rangle$ form an IRD-pattern of length n. Indeed, for any $f : n \to \omega$, the type

$$\{R(x; b_{i,j}) : i < n, j < f(i)\} \cup \{\neg R(x; b_{i,j}) : i < n, f(i) \leq j < \omega\}$$

is consistent by saturation.

If we assume that T is o-minimal, since dp-rank and op-dimension coincide, then this collapse of multi-order indiscernibles characterizes dp-rank.

4. Rosiness and indsicernible function spaces

4.1. **Background on rosiness.** There are several equivalent ways to define rosy theories, but for the sake of brevity, we present the one that is closest to the work we will carry in this article. This definition is based on the following notion of "equivalence relation ranks" as proposed in [2].

We remark that when defined in terms of local þ-ranks or the behavior of þ-independence, it is very important to work with a theory that eliminates imaginaries; that is, the class of rosy theories is strictly contained the class of "real-rosy theories." However, the characterization in terms of equivalence relation ranks eliminates this discrepancy – a theory T has finite equivalence relation ranks if and only if T^{eq} has finite equivalence relation ranks.

Definition 4.1 (Equivalence-relation-rank: $\mathsf{eq\text{-}rk}$)**.** Let $\pi(x)$ be a partial type, and let Δ be a finite set of formulas $\delta(x, x', z)$ such that for every c, $\delta(x, x', c)$ defines an equivalence relation. (We call such a Δ a *finite set of families of equivalence relations.*)

- $\mathsf{eq\text{-}rk}(\pi(x), \Delta) \geq 0$
- For a limit ordinal λ, $\mathsf{eq\text{-}rk}(\pi(x), \Delta) \geq \lambda$ if $\mathsf{eq\text{-}rk}(\pi(x), \Delta) \geq \alpha$ for all $\alpha < \lambda$.
- $\mathsf{eq\text{-}rk}(\pi(x), \Delta) \geq \alpha + 1$ if there are a formula $\delta(x, x', z) \in \Delta$ and a parameter c such that, with $E(x, x') = \delta(x, x', c)$, the set,

  $$\{b/E : b \vDash \pi \wedge \mathsf{eq\text{-}rk}(\pi(x) \cup \{E(x, b)\}, \Delta) \geq \alpha\}$$

  is infinite.
- $\mathsf{eq\text{-}rk}(\pi(x), \Delta) = \infty$ if $\mathsf{eq\text{-}rk}(\pi(x), \Delta) \geq \alpha$ for every ordinal α.

Remark 4.2. Let $\pi(x)$ be a partial type, and let Δ be a finite set of families of equivalence relations. By compactness, if $\mathsf{eq\text{-}rk}(\pi(x), \Delta) \geq \omega$, then $\mathsf{eq\text{-}rk}(\pi(x), \Delta) = \infty$.

The key theorem of [2] (for our purposes) is the following.

Theorem 4.3 ([2])**.** *Let T be a theory with big model $\mathbb{U}$. T is rosy if and only if $\mathsf{eq\text{-}rk}(\pi(x), \Delta) < \omega$ for every partial type $\pi(x)$ and every finite set Δ of families of equivalence relations.*

In the remainder of this subsection, we define a variant of the equivalence relation ranks as a coding property, and we show that non-rosiness can also be characterized by "coding equivalence relations."

Definition 4.4. Let T be a theory with big model $\mathbb{U}$. For a formula $\varphi(x; y_0, y_1)$, we say that φ *codes equivalence relations* if for every $0 < n < \omega$, there are $(a_i)_{i<n}$ of sort x and $(b_f)_{f \in {}^n n}$ of sort y such that for all $i < n$ and $f_0, f_1 \in {}^n n$,

$$\vDash \varphi(a_i; b_{f_0}, b_{f_1}) \;\Leftrightarrow\; f_0(i) = f_1(i).$$

Lemma 4.5. *Let T be a theory with big model $\mathbb{U}$, and assume T eliminates imaginaries. T is non-rosy if and only if there is a partitioned formula $\varphi(x; y_0, y_1)$ that codes equivalence relations in $\mathbb{U}$.*

Proof. First, suppose $\varphi(x; y_0, y_1)$ codes equivalence relations. Then, let

$$\delta(y_0, y_1; x) = \forall y \left(\varphi(x; y_0, y) \leftrightarrow \varphi(x; y_1, y)\right).$$

By compactness and the fact that $\varphi(x; y_0, y_1)$ codes equivalence relations, there are $(a_i)_{i<\omega}$ and $(b_f)_{f\in {}^{\omega}\omega}$ such that for all $i < \omega$ and $f_0, f_1 \in {}^{\omega}\omega$,

$$\vDash \varphi(a_i; b_{f_0}, b_{f_1}) \;\Leftrightarrow\; f_0(i) = f_1(i).$$

Now, we make a simple observation: Let $\sigma \in {}^{<\omega}\omega$, and for each $j < \omega$, let $\sigma^\frown j \subset f_j \in {}^{\omega}\omega$; then,

$$\mathsf{eq\text{-}rk}\left(\bigwedge_{\ell<|\sigma|} \delta(y, b_{f_0}; a_\ell) \wedge \delta(y, b_{f_j}; a_{|\sigma|}), \{\delta\}\right) \geq 0$$

for every $j < \omega$. From this, it follows immediately that $\mathsf{eq\text{-}rk}(x{=}x, \{\delta\}) \geq \omega$, and this shows that T is not rosy.

Conversely, suppose T is not rosy. Through simple coding tricks, we may assume that there is a single family of equivalence relations $\delta(x, x'; z)$ such that $\mathsf{eq\text{-}rk}(x{=}x, \{\delta\}) = \infty$. From this, we recover $(c_\tau)_{\tau\in {}^{<\omega}\omega}$ and $(d_\sigma)_{\sigma\in {}^{<\omega}\omega}$ such that, for any $\sigma \in {}^{<\omega}\omega$ and $\tau = \sigma^\frown j$ $(j < \omega)$

$$\mathsf{eq\text{-}rk}\left(\bigwedge_{\ell<|\sigma|} \delta(x, c_{\tau\restriction(\ell+1)}; d_{\sigma\restriction\ell}), \{\delta\}\right) \geq \omega.$$

By the Pigeonhole Principle, we may assume that $c_{\sigma^\frown j} d_\sigma \equiv c_{\sigma'^\frown j'} d_{\sigma'}$ whenever $|\sigma| = |\sigma'|$. Now, let $0 < n < \omega$ be given. Applying several automorphisms, we may assume that $d_\sigma = d_{\sigma'} =: a_{|\sigma|}$ whenever $\sigma, \sigma' \in {}^{<n}n$ satisfy $|\sigma| = |\sigma'|$; then, for $f \in {}^{n}n$, we choose any $b_f \vDash \bigwedge_{\ell<n} \delta(x, c_{f\restriction(\ell+1)}; d_{f\restriction\ell})$. As n was arbitrary, we have shown that $\varphi(x; y_0, y_1) = \delta(y_0, y_1; x)$ codes equivalence relations. □

4.2. **Ramsey theory for function spaces.** We have established that non-rosiness is equivalent to "coding equivalence relations." In this subsection, develop of theory of indiscernibles that corresponds to the latter notion.

Proposition 4.6. *$(T_{\mathrm{fn}}, \mathscr{L}_{\mathrm{fn}})$ is a theory of indiscernibles.*

Proof. First we prove that $(T_{\mathrm{eq}:n}, \mathscr{F}_{\mathrm{fn}})$ is a theory of indiscernibles. The base case follows since $(T_{\mathrm{eq}:1}, \mathscr{F}_{\mathrm{fn}})$ is interdefinable with the theory of linear order.

Now suppose the generic model of $T_{\mathrm{eq}:n}$ has extremely amenable automorphism group A, and consider $T_{\mathrm{eq}:(n+1)}$. Call the E_{x_i} simply E_i for $0 \leq i \leq n$. The X sort is finite and rigid. Thus the automorphism group acts on the $\bigcap_{i\leq n-1} \leq E_i$ classes (and permutes any point(s) in the classes.) Adding a new equivalence relation cross-cuts each $\bigcap_{i\leq n-1} \leq E_i$ class. Thus the automorphism group of the generic model of $T_{eq:(n+1)}$ is $\mathrm{Aut}(\mathbb{Q}, <)^{\mathbb{Q}} \rtimes A$, which we know to be extremely amenable.

To complete the argument for $(T_{\mathrm{fn}}, \mathscr{F}_{\mathrm{fn}})$, there are two ways to proceed. Firstly, one can observe that the automorphism groups of the generic models of the $(T_{\mathrm{eq}:n}, \mathscr{F}_{\mathrm{fn}})$'s form a directed family of subgroups of the automorphism group of the generic model of $(T_{\mathrm{fn}}, \mathscr{F}_{\mathrm{fn}})$, and that union of this family is dense the latter; by Lemma 6.7 of [5], it follows that $(T_{\mathrm{fn}}, \mathscr{F}_{\mathrm{fn}})$ is a theory of indiscernibles. A second approach would be to use compactness and the Modeling Properties of the $(T_{\mathrm{eq}:n}, \mathscr{F}_{\mathrm{fn}})$'s to prove that $(T_{\mathrm{fn}}, \mathscr{F}_{\mathrm{fn}})$ also has the Modeling Property. □

Now that we know what theory of indiscernibles to work with, we can state the main characterization theorem of this section. Its proof comprises all of the next subsection.

Theorem 4.7. *Let T be a theory with big model $\mathbb{U}$, and assume T eliminates imaginaries. T is rosy if and only if every indiscernible T_{fn}-picture in $\mathbb{U}$ is an indiscernible T_2-picture via $\mathsf{X}_0 = \mathsf{X}$, $\mathsf{X}_1 = \mathsf{F}$.*

4.3. **Proof of Theorem 4.7.** The proof of Theorem 4.7 amounts to the conjunction of Propositions 4.8 and 4.10, which we prove in turn. Throughout this subsection, assume that T eliminates imaginaries.

Proposition 4.8. *Let T be a theory with big model $\mathbb{U}$. If T is rosy, then every indiscernible T_{fn}-picture in $\mathbb{U}$ is an indiscernible T_2-picture via $\mathsf{X}_0 = \mathsf{X}$, $\mathsf{X}_1 = \mathsf{F}$.*

Before proving this proposition, we will need to establish the following lemma.

Lemma 4.9. *Let T be a theory with big model $\mathbb{U}$. If there is an indiscernible T_{fn}-picture in $\mathbb{U}$ that is not an indiscernible T_2-picture via $\mathsf{X}_0 = \mathsf{X}$, $\mathsf{X}_1 = \mathsf{F}$, then there is a formula $\varphi(x; y, y')$ that codes equivalence relations in $\mathbb{U}$.*

Proof. Let $I : \mathcal{A} \to \mathbb{U}$ be an indiscernible picture of the countable model $\mathcal{A} \vDash T^*_{\text{fn}}$, and suppose I is not an indiscernible T_2-picture. In particular, there are $0 < n_{\mathsf{X}}, n_{\mathsf{F}} < \omega$, $a_0 < \cdots < a_{n_{\mathsf{X}}-1}$ and $a'_0 < \cdots < a'_{n_{\mathsf{X}}-1}$ in $\mathsf{X}^{\mathcal{A}}$, $b_0 < \cdots < b_{n_{\mathsf{F}}-1}$ and $b'_0 < \cdots < b'_{n_{\mathsf{F}}-1}$ in $\mathsf{F}^{\mathcal{A}}$, and a formula $\varphi(x; y)$ in the language of T such that:

- $x = (x_0, ..., x_{n_{\mathsf{X}}-1})$ and $y = (y_0, ..., y_{n_{\mathsf{F}}-1})$
- $\text{qftp}^{\mathcal{A}}(a, b) \neq \text{qftp}^{\mathcal{A}}(a', b')$
- $\mathbb{U} \vDash \varphi(Ia, Ib)$ and $\mathbb{U} \vDash \neg\varphi(Ia', Ib')$

By model-completeness and the indiscernibility of I, we may assume that $a' = a$. We will also take the lemma as obvious in the case that $n_{\mathsf{F}} = 1$ (using $\varphi(x, y) \wedge \varphi(x, y')$ to code equivalence relations), so in all that follows, we assume $n_{\mathsf{F}} \geq 2$.

Claim. *We may assume that $p_b(a, y) = \text{qftp}^{\mathcal{A}}(a, b)$ and $p_{b'}(a, y) = \text{qftp}^{\mathcal{A}}(a, b')$ differ only in the following way: There are $i_0 < n_{\mathsf{X}}$ and $t_0 < n_{\mathsf{F}}$ such that for all $i < n_{\mathsf{X}}$, $s, t < n_{\mathsf{F}}$,*

$$\begin{aligned} i \neq i_0 \;&\Rightarrow\; \mathcal{A} \vDash (E(a_i, b_s, b_t) \leftrightarrow E(a_i, b'_s, b'_t)) \\ t_0 \notin \{s, t\} \;&\Rightarrow\; \mathcal{A} \vDash (E(a_i, b_s, b_t) \leftrightarrow E(a_i, b'_s, b'_t)), \\ |b_{t_0}/E_{a_{i_0}}| \geq 2 \;&\wedge\; |b'_{t_0}/E_{a_{i_0}}| = 1. \end{aligned}$$

Proof of claim. Let G be the set of all functions $g : n_{\mathsf{X}} \times n_{\mathsf{F}} \times n_{\mathsf{F}} \to 2$ satisfying the following conditions:

(1) For all $i < n_{\mathsf{X}}$ and $t_0, t_1, t_2 < n_{\mathsf{F}}$,

$$t_0 = t_1 \Rightarrow g(i, t_0, t_1) = 1,\ g(i, t_0, t_1) = g(i, t_1, t_0),$$

$$g(i, t_0, t_1) = 1 \wedge g(i, t_1, t_2) = 1 \Rightarrow g(i, t_0, t_2) = 1.$$

(2) For all $k < n_{\mathsf{X}}$ and $i_0 < \cdots < i_{k-1} < i_k < n_{\mathsf{X}}$, for all $t_0 < n_{\mathsf{F}}$,

$$\left\{ t < n_{\mathsf{X}} : \bigwedge_{j \leq k} g(i_j, t, t_0) = 1 \right\}$$

is a convex subset of

$$\left\{ t < n_{\mathsf{X}} : \bigwedge_{j < k} g(i_j, t, t_0) = 1 \right\}.$$

Then G is precisely the set of functions $n_{\mathsf{X}} \times n_{\mathsf{F}} \times n_{\mathsf{F}} \to 2$ that correspond naturally to quantifier-free-complete types $p(a, y)$ extending $\mathrm{qftp}^{\mathcal{A}}(a)$; we write $p \mapsto g_p$ and $g \mapsto p_g$ for the two directions of the bijection. Let's also name the function $h : n_{\mathsf{X}} \times n_{\mathsf{F}} \times n_{\mathsf{F}} \to 2$ given by,

$$h(i, t_0, t_1) = \begin{cases} 1 & \text{if } t_0 = t_1 \\ 0 & \text{otherwise.} \end{cases}$$

It's clear that $h \in G$. For any $g \in G$, we define $g^* : n_{\mathsf{X}} \times n_{\mathsf{F}} \times n_{\mathsf{F}} \to 2$ as follows:

- For $t < n_{\mathsf{F}}$, we say that t is *g-isolated* if for all $i < n_{\mathsf{X}}$, $g(i, t, t') = 1 \Rightarrow t = t'$.
- Set
$$t^* = \max \{t : t \text{ is not } g\text{-isolated}\},$$
$$i^* = \max \{i : (\exists t)\, t \neq t^* \wedge g(i, t, t^*) = 1\}.$$
- Finally, define g^* by

$$g^*(i, t_0, t_1) = \begin{cases} g(i, t_0, t_1) & \text{if } i \neq i^* \text{ or } t^* \notin \{t_0, t_1\} \\ 1 & \text{if } i = i^* \text{ and } t_0 = t_1 = t^* \\ 0 & \text{if } i = i^* \text{ and } \{t^*\} \subsetneq \{t_0, t_1\}. \end{cases}$$

It is not hard to verify that g^* remains in G, and it is equally clear that for each $g \in G$, there is a number $m < \omega$ such that if $g_0 = g$, and $g_{k+1} = g_k^*$ for each $k < m$, then $g_m = h$. Consequently, if g_{p_b} and $g_{p_{b'}}$ are the members of G corresponding to $p_b(a, y)$ and $p_{b'}(a, y)$, then there is a "path" $g_0 = g_{p_b}, g_1, ..., g_{N-1}, g_N = g_{p_{b'}}$ such that for each $i < N$, either $g_{i+1} = g_i^*$ or $g_i = g_{i+1}^*$. Since $\mathbb{U} \vDash \varphi(Ia, Ib)$ and $\mathbb{U} \vDash \neg\varphi(Ia, Ib')$, by indiscernibility, there is an $i < N$ such that

$$c \vDash p_{g_i}(a, y) \Rightarrow \mathbb{U} \vDash \varphi(Ia, Ic),\ \ c \vDash p_{g_{i+1}}(a, y) \Rightarrow \mathbb{U} \vDash \neg\varphi(Ia, Ic).$$

Then, replacing p_b with p_{g_i} and $p_{b'}$ with $p_{g_{i+1}}$ is sufficient for the claim. □

Observation. By indiscernibility and model-completeness, we may further assume that $b_j = b'_j$ whenever $j \in n_{\mathsf{F}} \setminus \{t_0\}$.

Finally, taking $s_0 \in W \setminus \{t_0\}$ and $\varphi^\circ(x_{i_0}; y_{s_0}, y_{t_0}) =$

$$\varphi(a_{0,...,i_0-1}, x_{i_0}, a_{i_0+1,...,n_{\mathsf{X}}-1}, b_{0,...,s_0-1}, y_{s_0}, b_{s_0+1,...,t_0-1}, y_{t_0}, b_{t_0+1,...,n_{\mathsf{F}}-1})$$

(when $s_0 < t_0$) it is not difficult to verify (following [10]) that $\varphi^\circ(x; y, y')$ codes equivalence relations. (When $s_0 > t_0$, we use a similar definition of φ°.) □

Proof of Proposition 4.8. Proving the contrapositive, if there is an indiscernible T_{fn}-picture in $\mathbb{U}$ that is not an indiscernible T_2-picture via $\mathsf{X}_0 = \mathsf{X}$, $\mathsf{X}_1 = \mathsf{F}$, then by Lemma 4.9, there is a formula $\varphi(x; y, y')$ that codes function spaces in $\mathbb{U}$. Then, by Lemma 4.5, the existence of a formula that codes equivalence relations in $\mathbb{U}$ implies that T is not rosy. □

Proposition 4.10. *Let T be a theory with big model $\mathbb{U}$. If every indiscernible T_{fn}-picture in $\mathbb{U}$ is an indiscernible T_2-picture via $\mathsf{X}_0 = \mathsf{X}$, $\mathsf{X}_1 = \mathsf{F}$, then T is rosy.*

Proof. For the contrapositive, suppose T is not rosy. By Lemma 4.5, there is a formula $\varphi(x; y, y')$ that codes function spaces in $\mathbb{U}$. We will use this fact to construct an indiscernible T_{fn}-picture in $\mathbb{U}$ that is not an indiscernible T_2-picture via $\mathsf{X}_0 = \mathsf{X}$, $\mathsf{X}_1 = \mathsf{F}$.

In this paragraph, we collect the data necessary to devise an EM-template. By compactness, there are $(c_i)_{i<\omega}$ and $(d_f)_{f \in {}^\omega\omega}$ in $\mathbb{U}$ such that

$$\vDash \varphi(c_i; d_{f_0}, d_{f_1}) \Leftrightarrow f_0(i) = f_1(i)$$

for all $i < \omega$ and $f_0, f_1 : \omega \to \omega$. Let $\mathcal{A} \vDash T^*_{\text{fn}}$ be the generic model, and let $u : \mathsf{X}^{\mathcal{A}} \to \omega$ and $v : \mathsf{F}^{\mathcal{A}} \to {}^{\omega}\omega$ be injective functions such that

$$\mathcal{A} \vDash E(a, b_0, b_1) \Leftrightarrow \mathbb{U} \vDash \varphi(c_{u(a)}; d_{v(b_0)}, d_{v(b_1)})$$

for all $a \in \mathsf{X}^{\mathcal{A}}$ and $b_0, b_1 \in \mathsf{F}^{\mathcal{A}}$. From these, we define our EM-template by setting

$$\Gamma(\overline{a}^\frown \overline{a}') = (c_{u(a_i)})_{i<m}{}^\frown (d_{v(a'_j)})_{j<n}$$

whenever $\overline{a} \in (\mathsf{X}^{\mathcal{A}})^m$ and $\overline{a}' \in (\mathsf{F}^{\mathcal{A}})^n$.

By the Modeling Property of $(T_{\text{fn}}, \mathscr{F}_{\text{fn}})$, we recover an indiscernible picture $I : \mathcal{A} \to \mathbb{U}$ patterned on Γ, and we take it as obvious that I is not an indiscernible T_2-picture via $\mathsf{X}_0 = \mathsf{X}$, $\mathsf{X}_1 = \mathsf{F}$. □

5. NTP2 and Convexly-ordered Equivalence Relations

Definition 5.1 (k-TP2)**.** Let T be a theory with big model $\mathbb{U}$, and let $\varphi(x; y)$ be a partitioned formula of the language $\mathscr{L}$ of T. For $2 \leq k < \omega$, $\varphi(x; y)$ is said to have the *k-tree property of the second kind* (k-TP2) in T if there is a family $(b^i_j)_{i,j\in\omega}$ in $\mathbb{U}$ such that:

- For every function $f : \omega \to \omega$, $\left\{\varphi(x, b^i_{f(i)})\right\}_{i<\omega}$ is consistent.
- For every $i < \omega$, $\left\{\varphi(x, b^i_j)\right\}_{j<\omega}$ is k-inconsistent.

T is said to have TP2 if there are a formula $\varphi(x; y)$ and $2 \leq k < \omega$ such that φ has k-TP2 in T; otherwise, we say that T is NTP2.

Fact 5.2. (Theorem 0.2 of [12]) The following are equivalent:

- T has TP2.
- Some formula $\varphi(x; y)$ has 2-TP2 in T.

5.1. **ceq-Indiscernibles.** We will argue independently for why T^*_{ceq} is a theory of indiscernibles. By [5], the finite models of T_{ceq} form a Ramsey class, and so by [10], T^*_{ceq} is a theory of indiscernibles. These indiscernibles have already been widely used in the TP2 context, typically as *array indiscernibles* (as described in [7]) where the language is composed of two linear orderings.

5.2. **Characterizing NTP2.** We follow Claim 2.19 in [13] in the sense that the absence of TP2 in a theory is characterized by dynamics in the EM-type of indiscernible sequences. Namely, if the indiscernible sequence allows for a certain kind of consistency, it must then allow for more. Since we are tracking the dynamics of formulas asserting consistency, as they show up in the EM-type, it is helpful to make use of a notation from [9]:

Definition 5.3. The characteristic sequence associated to a formula $\varphi(x; y)$ is defined to be $P_n(y_0, \ldots, y_{n-1}) \equiv \exists x(\bigwedge_{i<n} \varphi(x; y_i))$

We will write $P^\varphi_n(y_0, \ldots, y_{n-1})$ to denote dependence on φ.

Definition 5.4 (Coding equivalence)**.** Let $\mathcal{A} \vDash T^*_{\text{ceq}}$ be a countable model, and let $I : \mathcal{A} \to \mathbb{U}^y$ be an indiscernible picture. For a formula $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$, we say that *$\varphi$ codes equivalence on I* if, for all $a_0, a_1 \in A$ with $a_0 <^{\mathcal{A}} a_1$,

$$E^{\mathcal{A}}(a_0, a_1) \text{ if and only if } \vDash \neg\varphi(Ia_0; Ia_1).$$

Lemma 5.5. *Let T be a theory with big model $\mathbb{U}$. Let $I : \mathcal{A} \to \mathbb{U}$ be an indiscernible picture of $\mathcal{A} \vDash T^*_{ceq}$. Then I does not collapse to order if and only if there exist a formula $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$ and an embedding $f : \mathcal{A} \to \mathcal{A}$ such that φ codes equivalence on $I \circ f$.*

Proof. As before, one direction is clear, so we show the other direction. Suppose I does not collapse to order. Then, there are $<^{\mathcal{A}}$-increasing n-tuples $\overline{a}_1, \overline{a}_2$ from I such that $\text{qftp}^{\mathcal{A}}(\overline{a}_1) = \text{qftp}^{\mathcal{A}}(\overline{a}_2)$ but $I\overline{a}_1, I\overline{a}_2$ do not have the same complete type in $\mathbb{U}$. Thus for some formula $\theta(\overline{x})$ of the language of $\mathscr{L}$ of T, $\vDash \theta(I\overline{a}_1)$ and $\vDash \neg\theta(I\overline{a}_2)$. Let $q_i = \text{qftp}^{\mathcal{A}}(\overline{a}_i)$, $i = 1, 2$. In general, let $S_n^{\text{ceq}}(\emptyset)$ be the set of all quantifier-free n-types $p(x_0, ..., x_{n-1})$ where

- $p \vdash x_i < x_j$ for all $i < j < n$,
- $p \vdash E(x_i, x_i)$ for all $i < n$,
- if $p \vdash E(x_i, x_j)$, then $p \vdash E(x_j, x_i)$ for all $i < j < n$,
- if $p \vdash E(x_i, x_j) \wedge E(x_j, x_k)$ for some $i < j < k < n$, then $p \vdash E(x_i, x_k)$, and
- if $p \vdash E(x_i, x_j)$ for $i < k < j < n$, then $p \vdash E(x_i, x_k) \wedge E(x_k, x_j)$.

That is, $S_n^{\text{ceq}}(\emptyset)$ are the set of all quantifier-free n-types realized by some $<^{\mathcal{A}}$-ascending sequence of n-tuples in $\mathcal{A}$. Hence $q_1, q_2 \in S_n^{\text{ceq}}(\emptyset)$. Now form a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = S_n^{\text{ceq}}(\emptyset)$ and, for $p_1, p_2 \in S_n^{\text{ceq}}(\emptyset)$, $\{p_1, p_2\} \in \mathcal{E}$ if there exists $0 < j < n$ such that

- $p_1(x_0, ..., x_{j-1}, x_{j+1}, ..., x_{n-1}) = p_2(x_0, ..., x_{j-1}, x_{j+1}, ..., x_{n-1})$,
- $p_1 \vdash \neg E(x_{j-1}, x_{j+1})$ (hence $p_2 \vdash \neg E(x_{j-1}, x_{j+1})$),
- $p_1 \vdash E(x_{j-1}, x_j)$ or $p_1 \vdash E(x_j, x_{j+1})$, and
- $p_2 \vdash \neg E(x_{j-1}, x_j) \wedge \neg E(x_j, x_{j+1})$.

That is, p_1 thinks x_j is E-related to either x_{j-1} or x_{j+1} (but not both, as they are E-unrelated) and p_2 thinks that x_j is E-unrelated to x_{j-1} and x_{j+1}. To each edge $\{p_1, p_2\} \in \mathcal{E}$, we can associate a number $1 \leq m < n$ of the size of the E-class in p_2 to which p_1 thinks x_j is E-related, namely $m := \max\{m' : m' \leq j, p_2 \vdash E(x_{j-m'}, x_{j-1})\}$ if $p_1 \vdash E(x_{j-1}, x_j)$ and $m := \max\{m' : m' < n-j, p_2 \vdash E(x_{j+1}, x_{j+m'})\}$ if $p_1 \vdash E(x_j, x_{j+1})$. Call this $m(\{p_1, p_2\}) = m$. The reader should check that $\mathcal{G}$ is connected (one can get between any two configurations by pulling apart or adding on a single element to an E-class iteratively). Since I is an indiscernible picture, it induces a vertex 2-coloring of $\mathcal{G}$, namely $\eta : \mathcal{V} \to \{0, 1\}$ where $\eta(p) = 1$ if and only if, for some (equivalently any) realization $\overline{a} \vDash p$ (with $\overline{a} \in A^n$), $\vDash \theta(I\overline{a})$. Hence, $\eta(q_1) = 1$ and $\eta(q_2) = 0$. In particular, η is non-constant. Now choose an edge $\{p_1, p_2\} \in \mathcal{E}$ such that

- $\eta(p_1) \neq \eta(p_2)$, and
- $m(\{p_1, p_2\})$ is minimal such.

Since $\mathcal{G}$ is connected and η is non-constant, such an edge exists. Let j witness that $\{p_1, p_2\}$ is an edge, without loss suppose $p_1 \vdash E(x_{j-1}, x_j)$ (the argument for $p_1 \vdash E(x_j, x_{j+1})$ is similar), and let $m = m(\{p_1, p_2\})$. Without loss of generality, suppose $\eta(p_1) = 0$ and $\eta(p_2) = 1$ (by possibly swapping θ and $\neg\theta$). Notice that, by minimal choice of $m(\{p_1, p_2\})$, we have in fact that, for all $p \in S_n^{\text{ceq}}(\emptyset)$, if

$$p(x_0, ..., x_{j-m-1}, x_{j-1}, ..., x_{n-1}) = p_2(x_0, ..., x_{j-m-1}, x_{j-1}, ..., x_{n-1}), \tag{1}$$

then $\eta(p) = \eta(p_2) = 1$ (otherwise, we could find a path from p to p_2 maintaining this condition of equality, hence contradicting the minimality of m). Similarly, for all $p \in S_n^{\text{ceq}}(\emptyset)$, if

$$p(x_0, ..., x_{j-m-1}, x_{j-1}, ..., x_{n-1}) = p_1(x_0, ..., x_{j-m-1}, x_{j-1}, ..., x_{n-1}), \tag{2}$$

then $\eta(p) = \eta(p_1) = 0$. Now, for some realization $\overline{a} \vDash p_1$, let

$$\varphi(x; y) := \theta(Ia_0, ..., Ia_{j-2}, x, y, Ia_{j+1}, ..., Ia_{n-1}).$$

By ultrahomogeneity, there exists an embedding $f : \mathcal{A} \to \mathcal{A}$ such that $a_{j-2} <^{\mathcal{A}} f(A) <^{\mathcal{A}} Ia_{j+1}$ and each element of $f(A)$ is $E^{\mathcal{A}}$ unrelated to both a_{j-2} and a_{j+1}. We claim φ codes equivalence on $I \circ f$, as desired.

Fix $b_0 <^{\mathcal{A}} b_1$ from $f(A)$ and let $p(\overline{x}) = \text{qftp}^{\mathcal{A}}(a_0, ..., a_{j-2}, b_0, b_1, a_{j+1}, ..., a_{n-1})$. In particular, since $\neg E^{\mathcal{A}}(b_1, a_{j+1})$ and $m \geq 1$, one of the following holds:

- $E^{\mathcal{A}}(b_0, b_1)$, hence (2) holds, hence $\eta(p) = 1$, hence $\vDash \neg\varphi(Ib_0; Ib_1)$; or
- $\neg E^{\mathcal{A}}(b_0, b_1)$, hence (1) holds, hence $\eta(p) = 0$, hence $\vDash \varphi(Ib_0; Ib_1)$.

Therefore, $\vDash \varphi(Ib_0; Ib_1)$ if and only if $\neg E^{\mathcal{A}}(b_0, b_1)$, as desired. □

Recall the general definition of dividing across, given in Definition 2.23. For the specific case of convex equivalence relation indiscernibles, we define dividing vertically. One could also define the corresponding notion of dividing laterally.

Definition 5.6 (Dividing vertically)**.** Let $\mathcal{A} \vDash T^*_{\text{ceq}}$ be a countable model, and let $I : \mathcal{A} \to \mathbb{U}^y$ be an indiscernible picture. For a formula $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$, we say that φ *divides vertically across* I if there exists an indiscernible sequence $\langle a_i \rangle_{i<\omega}$ on $\mathcal{A}$ with $\neg E^{\mathcal{A}}(a_0, a_1)$ such that

$$\{\varphi(x; Ia_i)\}_{i<\omega} \text{ is inconsistent.}$$

Remark 5.7 (Coding equivalence implies does not divide vertically). Fix $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$. If φ codes equivalence on I, then φ does not divide vertically across I. This is because, for any finite set $B_0 \subseteq A$ of pairwise $E^{\mathcal{A}}$-unrelated elements, there exists an element $a \in A^n$ such that $a <^{\mathcal{A}} b$ and $\neg E^{\mathcal{A}}(a, b)$ for all $b \in B_0$. Therefore, $\vDash \varphi(Ia; Ib)$ for all $b \in B_0$. By compactness, for any $B \subseteq A$ of pairwise $E^{\mathcal{A}}$-unrelated elements, $\{\varphi(x; Ib)\}_{b\in B}$ is consistent.

Theorem 5.8. *Let T be a theory with big model $\mathbb{U}$. The following are equivalent:*

(1) *T is NTP2.*
(2) *For every indiscernible T_{ceq}-picture $I : \mathcal{A} \to \mathbb{U}$, for every formula $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$, if $\varphi(x, y)$ does not divide vertically across I, then $\varphi(x; y)$ does not divide across I.*

Proof of 1⇒2. Let $I : \mathcal{A} \to \mathbb{U}$ be an indiscernible picture of $\mathcal{A} \vDash T^*_{\text{ceq}}$ and suppose $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$ is such that φ does not divide vertically across I. Suppose φ divides across I – that is, we have an indiscernible sequence $\langle a_i \rangle_{i<\omega}$ in $\mathcal{A}$ such that $\{\varphi(x; Ia_i)\}_{i<\omega}$ is inconsistent. Without loss, suppose $a_i <^{\mathcal{A}} a_j$ for all $i < j < \omega$. The only two possibilities for $\langle a_i \rangle_{i<\omega}$ are then

- $E^{\mathcal{A}}(a_i, a_j)$ for all $i < j < \omega$, or
- $\neg E^{\mathcal{A}}(a_i, a_j)$ for all $i < j < \omega$.

Since φ does not divide vertically across I, the second case cannot hold. Therefore, the first case holds. However, by compactness, $\{\varphi(x; Ia_i)\}_{i<k}$ is inconsistent for some $k < \omega$. Since I is an indiscernible picture, for any $E^{\mathcal{A}}$-related $B_0 \subseteq A$ with $|B_0| = k$, $\{\varphi(x; Ib)\}_{b\in B_0}$ is inconsistent. Therefore, for any $A' \subseteq A$ so that $(A'; \lhd)$ is isomorphic to $(\omega \times \omega, <_{\text{lex}}, \sim_1)$ (where $(i_0, j_0) \sim_1 (i_1, j_1)$ iff $i_0 = i_1$), we have that $\langle Ia \rangle_{a\in A'}$ is a witness to the fact that $\varphi(x; y)$ has k-TP2 in T. Therefore, T has TP2. This completes the proof of 1⇒2. □

Proof of 2⇒1. Suppose T has TP2. Let $\varphi(x; y)$ witness 2-TP2 in T, and let $\mathcal{A} \vDash T^*_{\text{ceq}}$ be the generic model. Thus there are parameters $B := \{a^i_j : i, j < \omega\}$ in $\mathbb{U}$ such that:

- For any finite partial function $f : \omega \to \omega$ with domain $\{i_0 < \cdots < i_{n-1}\}$,

$$\vDash P^{\varphi}_n \left(b^{i_0}_{f(i_0)}, ..., b^{i_{n-1}}_{f(i_{n-1})} \right);$$

- For all $i, j_0, j_1 < \omega$, $\vDash \neg P^{\varphi}_2(b^i_{j_0}, b^i_{j_1})$.

Let $\mathscr{L}^+$ be the expansion of $\mathscr{L}$ (the language of T) with one new sort for $\mathcal{A}$ and a new function symbol $h : A \to \mathbb{U}_y$, and let Γ be the set of sentences of $\mathscr{L}^+(A)$ that includes the diagram $diag(\mathcal{A})$ of $\mathcal{A}$, the elementary diagram $ediag(\mathcal{M})$ of a model $\mathcal{M} \prec \mathbb{U}$ containing B, and the sentences,

$$E^{\mathcal{A}}(a, a') \to \neg P^{\varphi}_2(h(a), h(a')), \quad \bigwedge_{j<k<n} \neg E^{\mathcal{A}}(a_j, a_k) \to P^{\varphi}_n(h(a_0), ..., h(a_{n-1}))$$

for all $a, a', a_0, ..., a_{n-1} \in A$. Using B to verify finite-satisfiability, it's clear that Γ has a model, and from this model, we recover a picture $J : \mathcal{A} \to \mathbb{U}$ (as a restriction of the interpretation of h) such that

$$E^{\mathcal{A}}(a, a') \Rightarrow \neg P_2^{\varphi}(J(a), J(a')) \tag{3}$$

$$\bigwedge_{j<k<n} \neg E^{\mathcal{A}}(a_j, a_k) \Rightarrow P_n^{\varphi}(J(a_0), ..., J(a_{n-1})) \tag{4}$$

for all $a, a', a_0, ..., a_{n-1} \in A$. Using J to make an EM-template $A^{<\omega} \to \mathbb{U}^{<\omega} : \overline{a} \mapsto J\overline{a}$, and with the Modeling Property of T_{ceq}, we obatin an indiscernible picture $I : \mathcal{A} \to \mathbb{U}_y$ patterned on $\overline{a} \mapsto J\overline{a}$ – so also satisfying statements (3,4) with I in place of J. In particular,

$$\mathcal{A} \vDash a_0 < a_1 \wedge E(a_0, a_1) \ \Rightarrow \vDash \neg P_2^{\varphi}(I(a_0), I(a_1))$$
$$\mathcal{A} \vDash a_0 < a_1 \wedge \neg E(a_0, a_1) \ \Rightarrow \vDash P_2^{\varphi}(I(a_0), I(a_1))$$

Thus the indiscernible picture given by I does not collapse to order.

Moreover, by compactness and (3), for any set of pairwise $E^{\mathcal{A}}$-unrelated elements $B \subseteq A$, $\{\varphi(x; Ib)\}_{b \in B}$ is consistent. Thus, φ does not divide vertically across I. On the other hand, for any set of pairwise $E^{\mathcal{A}}$-related elements $B \subseteq A$, $\{\varphi(x; Ib)\}_{b \in B}$ is inconsistent. Hence φ divides across I. □

Therefore, we get a dichotomy for theories with NTP2: Either I collapses to order or there exists $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$ and an embedding $f : \mathcal{A} \to \mathcal{A}$ such that φ codes equivalence on $I \circ f$. Hence φ does not divide vertically across $I \circ f$. Then, the theorem says that φ does not divide across $I \circ f$. In summary:

Corollary 5.9. *Let T be a theory with big model $\mathbb{U}$. The following are equivalent:*

1. *T is NTP2.*
2. *For every indiscernible T_{ceq}-picture $I : \mathcal{A} \to \mathbb{U}$, either I is an indiscernible T_1-picture or, for every formula $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$ that does not divide vertically across I, $\varphi(x; y)$ does not divide across I.*

Observation. Note that the first condition in (2) above that I be an indiscernible T_1-picture is a special case of the second condition, that for every formula $\varphi(x; y) \in \mathscr{L}(\mathbb{U})$ that does not divide vertically across I, $\varphi(x; y)$ does not divide across I. This is because if I collapses to order, then the type of $I(\overline{a})$ in $\mathbb{U}$ is constant for $<^{\mathcal{A}}$-increasing tuples $\overline{a}$ from $\mathcal{A}$, and so there is no distinction between dividing vertically across I and dividing across I.

Also note that the first condition in (2) on its own would not be good enough to imply (1) from the Corollary; look at the case where $\mathbb{U}$ is a big model of T^*_{ceq}. We can argue that the theory is NIP, because basic relations are NIP, and therefore the theory is NTP2. However, it will not be the case that every indiscernible T^*_{ceq}-picture is an indiscernible T_1-picture. Consider the example where the identity $J : \mathcal{A} \to \mathcal{A}$ is used to make the EM-template on which indiscernible $I : \mathcal{A} \to \mathbb{U}$ is patterned. Clearly I will not be an indiscernible T_1-picture, as the equivalence relation is present.

One way to think of the (1) ⇒ (2) direction of the above Corollary is that, in a TP2 theory, the specific formula witnessing non-collapse to order may not express a distinction between consistency and inconsistency for φ-types, for some φ, as that would create a witness to TP2. Thus we have a partial notion of collapse: In TP2 theories and for indiscernible T_{ceq}-pictures I, there is some restricted set of formulas $\theta(\overline{x}) \in \mathscr{L}(\mathbb{U})$ such that if $<^{\mathcal{A}}$-increasing $\overline{a}, \overline{b}$ give $I(\overline{a}), I(\overline{b})$ disagreeing on their complete type in $\mathbb{U}$, it must be disagreement on some θ from this restricted set.

References

[1] M. Bodirsky, *New Ramsey classes from old*, Electron. J. Combin. **21** (2014), no. 2, Paper 2.22, 13.

[2] C. Ealy and A. Onshuus, *Characterizing rosy theories*, J. Symbolic Logic **72** (2007), no. 3, 721-1080.

[3] V. Guingona and C. D. Hill, *On a common generalization of Shelah's 2-rank, dp-rank, and o-minimal dimension*, Ann. Pure Appl. Logic. to appear.

[4] C. D. Hill, *A survey of generalized indiscernibles as model-complete theories* (2012). Preprint.

[5] A. S. Kechris, V. G. Pestov, and S. Todorcevic, *Fraïssé Limits, Ramsey Theory, and topological dynamics of automorphism groups*, Geometric and Functional Analysis GAFA **15**, no. 1, 106-189.

[6] B. Kim and H.-J. Kim, *Notions around tree property 1*, Annals of Pure and Applied Logic **162** (2011), 698-709.

[7] B. Kim, H.-J. Kim, and L. Scow, *Tree indiscernibilities, revisited*, Archive for Math. Logic **53** (2014), 211-232.

[8] K. Leeb, *Vorlesungen uber Pascaltheorie*, Lecture notes, Universitat Erlangen (1973).

[9] M. Malliaris, *The characteristic sequence of a first-order formula*, Journal of Symbolic Logic **75** (2010), no. 4, 1415-1440.

[10] L. Scow, *Characterization of NIP theories by ordered graph-indiscernibles*, Annals of Pure and Applied Logic **163** (2012), 1624-1641.

[11] ———, *Indiscernibles, EM-type, and Ramsey classes of trees*. preprint.

[12] S. Shelah, *Simple unstable theories*, Annals of Mathematical Logic **19** (1980), 177-203.

[13] ———, *Toward classifying unstable theories*, Annals Pure and Applied Logic **80** (1996), 229-255.

[14] ———, *Classification theory and the number of non-isomorphic models*, North-Holland Publishing Company, 1978.

[15] M. Sokić, *Ramsey property, ultrametric spaces, finite posets, and universal minimal flows*, Israel J. Math. **194** (2013), no. 2, 609–640.

[16] ———, *Semilattices and Ramsey property*. preprint.